

Nested Weighted Limit-Average Automata of Bounded Width

Krishnendu Chatterjee¹, Thomas A. Henzinger¹, and Jan Otop²

- 1 IST Austria
 {krish.chat,tah}@ist.ac.at
 2 University of Wrocław
 jotop@cs.uni.wroc.pl

Abstract

While weighted automata provide a natural framework to express quantitative properties, many basic properties like average response time cannot be expressed with weighted automata. Nested weighted automata extend weighted automata and consist of a master automaton and a set of slave automata that are invoked by the master automaton. Nested weighted automata are strictly more expressive than weighted automata (e.g., average response time can be expressed with nested weighted automata), but the basic decision questions have higher complexity (e.g., for deterministic automata, the emptiness question for nested weighted automata is PSPACE-hard, whereas the corresponding complexity for weighted automata is PTIME). We consider a natural subclass of nested weighted automata where at any point at most a bounded number k of slave automata can be active. We focus on automata whose master value function is the limit average. We show that these nested weighted automata with bounded width are strictly more expressive than weighted automata (e.g., average response time with no overlapping requests can be expressed with bound $k = 1$, but not with non-nested weighted automata). We show that the complexity of the basic decision problems (i.e., emptiness and universality) for the subclass with k constant matches the complexity for weighted automata. Moreover, when k is part of the input given in unary we establish PSPACE-completeness.

1 Introduction

Traditional to quantitative verification. In contrast to the traditional view of formal verification that focuses on Boolean properties of systems, such as “every request is eventually granted”, quantitative specifications consider properties like “the long-run average success rate of an operation is at least one half” or “the long-run average response time is below a threshold.” Such properties are crucial for performance related properties, for resource-constrained systems, such as embedded systems, and significant attention has been devoted to them [20, 14, 13, 21, 2].

Weighted automata. A classical model to express quantitative properties is *weighted automata* that extends finite automata where every transition is assigned a rational number called a *weight*. Each run results in a sequence of weights, and a *value function* aggregates the sequence into a single value. For non-deterministic weighted automata, the value of a word is the infimum value of all runs over the word. Weighted automata provide a natural and flexible framework to express quantitative¹ properties [14]. Weighted automata have been studied over finite words with weights from a semiring [20], and extended to infinite

¹ We use the term “quantitative” in a non-probabilistic sense, which assigns a quantitative value to each infinite run of a system, representing long-run average or maximal response time, or power consumption, or the like, rather than taking a probabilistic average over different runs.



licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

words with limit averaging or supremum as a value function [14, 13, 12]. While weighted automata over semirings can express several quantitative properties [26], they cannot express long-run average properties that weighted automata with limit averaging can [14]. However, even weighted automata with limit averaging cannot express the basic quantitative property of average response time [16, Example 5].

Nested weighted automata. To express properties like average response time, weighted automata were extended to *nested weighted automata* (NWA) [16]. An NWA consists of a master automaton and a set of slave automata. The master automaton runs over infinite input words. At every transition the master automaton can invoke a slave automaton that runs over a finite subword of the infinite word, starting at the position where the slave automaton is invoked. Each slave automaton terminates after a finite number of steps and returns a value to the master automaton. Each slave automaton is equipped with a value function for finite words, and the master automaton aggregates the returned values from slave automata using a value function for infinite words. For Boolean finite automata, nested automata are as expressive as the non-nested counterpart, whereas NWA are strictly more expressive than non-nested weighted automata [16]. It has been shown in [16] that NWA provide a specification framework where many basic quantitative properties, which cannot be expressed by weighted automata, can be expressed easily, and it provides a natural framework to study quantitative run-time verification.

The basic decision questions. We consider the basic automata-theoretic decision questions of emptiness and universality. The importance of these basic questions in the weighted automata setting are as follows: (1) Consider a system modeled by a finite-automata recognizing traces of the system and a quantitative property given as a weighted automaton or NWA. Then whether the worst-case (resp., best-case) behavior has the value at least λ is the emptiness (resp., universality) question on the product. (2) Problems related to model measuring (that generalizes model checking) and model repair also reduces to the emptiness problem [24, 16].

Complexity gap. In this work we focus on the following classical value functions: LIMAVG for infinite words, which is the long-run average property; and SUM, SUM⁺ (where SUM⁺ is the sum of absolute values) for finite words. While NWA are strictly more expressive than weighted automata, the complexity of the decision questions are either unknown or considerably higher. Table 1 (non-bold-faced results) summarizes the existing results for weighted automata [14] and NWA [16], for example, for NWA for SUM⁺ the known bounds are EXPSpace and PSPACE-hard, and for SUM even the decidability of the basic decision questions is open (or undecidable). Thus, a fundamental question is whether there exist sub-classes of NWA that are strictly more expressive than weighted automata and yet have better complexity than general NWA. We address this question in this paper.

Nested weighted automata with bounded width. For NWA, let the maximum number of slave automata that can be active at any point be the *width* of the automaton. In this work we consider a natural special class of NWA, namely, NWA with bounded width, i.e., NWA where at any point at most k slave automata can be active. For example, the average response time with bounded number of requests pending at any point can be expressed as NWA with bounded width, but not with weighted automata. Moreover, the class of NWA with bounded width is equivalent to automata with monitor counters [17], which are automata equipped with counters, where at each transition, a counter can be started, terminated, or the value of the counter can be increased or decreased. The transitions do not depend on the counter values, and hence they are referred to as monitor counters. The values of the counters when they are terminated gives rise to the sequence of weights, which is aggregated into a

	Deterministic (Emptiness/Universality)	Nondeterministic Emptiness	Nondeterministic Universality
Weighted aut.	PTIME		Undecidable
NWA (LIMAVG, SUM ⁺)	EXPSpace, PSPACE-hard PTIME (width k is constant) PSpace-c. (bounded width)		Undecidable
NWA (LIMAVG, SUM)	Open PTIME (width k is constant) PSpace-c. (bounded width)		Undecidable

■ **Table 1** Decidability and complexity of emptiness and universality for weighted and nested weighted automata with LIMAVG value function and SUM and SUM⁺ value function for slave automata. Our results are bold faced. Moreover all PTIME results become NLOGSPACE-complete when the weights are specified in unary.

single value with the LIMAVG value function (see [17]). Automata with monitor counters are similar in spirit with the class of register automata of [2].

Our contributions. Our contributions are as follows (summarized as bold-faced results in Table 1):

1. *Constant width.* We show that the emptiness problem (resp., the emptiness and the universality problems) for non-deterministic (resp., deterministic) NWA with constant width (i.e., k is constant) can be solved in polynomial time and is NLOGSPACE-complete when the weights are specified in unary. Thus we achieve the same complexity as weighted automata for a much more expressive class of quantitative properties.
2. *Bounded width.* We show that the emptiness problem (resp., the emptiness and the universality problems) for non-deterministic (resp., deterministic) NWA with bounded width (i.e., k is part of input given in unary) is PSPACE-complete. Thus we establish precise complexity when k is part of input given in unary.
3. *Deciding width.* We show that checking whether a given NWA has width k can be solved in polynomial time for constant k and in PSPACE if k is given in the input (Theorem 6).

Technical contributions. Our main technical contributions for deterministic (LIMAVG; SUM)-automata are as follows.

1. *Infinite infimum.* We first identify a necessary and sufficient condition for the infimum value over all words to be $-\infty$, and show that this condition can be checked efficiently.
2. *Lasso-approximation.* We show that if the above condition does not hold, then the infimum over all words can be approximated by lasso words, i.e., words of the form vu^ω . Moreover, we show that the infimum value is achieved with words where the slave automata runs for short length relative to the point of the invocation, and hence the partial averages converge.
3. *Reduction to width 1.* Using the lasso-approximation we reduce the emptiness problem of width bounded by k to the corresponding problem of width 1. We show that the case of width 1 can be solved using standard techniques.

Related works. Weighted automata over finite words have been extensively studied, the book [20] provides an excellent collection of results. Weighted automata on infinite words have been studied in [14, 13, 21]. The extension to weighted automata with monitor counters over finite words has been considered as cost register automata in [2]. A version of nested weighted automata over finite words has been studied in [6], and nested weighted automata over infinite words has been studied in [16]. Several quantitative logics have also been studied, such as [5, 7, 1]. In this work we consider a subclass of nested weighted automata which is strictly more expressive than weighted automata yet achieve the same complexity for the basic decision questions. Probabilistic models (such as Markov decision processes) with

quantitative properties (such as limit-average or discounted-sum) have also been extensively studied for single objectives [22, 27], and for multiple objectives and their combinations [19, 10, 15, 8, 18, 9, 23, 11, 3, 4]. While NWA with bounded width have been studied under probabilistic semantics [17], the basic automata theoretic decision problems have not been studied for them.

2 Preliminaries

2.1 Words and automata

Words. We consider a finite *alphabet* of letters Σ . A *word* over Σ is a (finite or infinite) sequence of letters from Σ . We denote the i -th letter of a word w by $w[i]$, and for $i < j$ we have that $w[i, j]$ is the word $w[i]w[i+1] \dots w[j]$. The length of a finite word w is denoted by $|w|$; and the length of an infinite word w is $|w| = \infty$. For an infinite word w , thus $w[i, \infty]$ is the suffix of the word with first $i - 1$ letters removed.

Labeled automata. For a set X , an X -labeled automaton \mathcal{A} is a tuple $\langle \Sigma, Q, Q_0, \delta, F, C \rangle$, where (1) Σ is the alphabet, (2) Q is a finite set of states, (3) $Q_0 \subseteq Q$ is the set of initial states, (4) $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation, (5) F is a set of accepting states, and (6) $C : \delta \mapsto X$ is a labeling function. A labeled automaton $\langle \Sigma, Q, q_0, \delta, F, C \rangle$ is *deterministic* if and only if δ is a function from $Q \times \Sigma$ into Q and Q_0 is a singleton. In definitions of deterministic labeled automata we omit curly brackets in the description of Q_0 and write $\langle \Sigma, Q, q_0, \delta, F, C \rangle$.

Semantics of (labeled) automata. A *run* π of a (labeled) automaton \mathcal{A} on a word w is a sequence of states of \mathcal{A} of length $|w| + 1$ such that $\pi[0]$ belong to the initial states of \mathcal{A} and for every $0 \leq i \leq |w| - 1$ we have $(\pi[i], w[i], \pi[i+1])$ is a transition of \mathcal{A} . A run π on a finite word w is *accepting* iff the last state $\pi[|w|]$ of the run is an accepting state of \mathcal{A} . A run π on an infinite word w is *accepting* iff some accepting state of \mathcal{A} occurs infinitely often in π . For an automaton \mathcal{A} and a word w , we define $\text{Acc}(w)$ as the set of accepting runs on w . Note that for deterministic automata, every word w has at most one accepting run ($|\text{Acc}(w)| \leq 1$).

Weighted automata. A *weighted automaton* is a \mathbb{Z} -labeled automaton, where \mathbb{Z} is the set of integers. The labels are called *weights*.

Semantics of weighted automata. We define the semantics of weighted automata in two steps. First, we define the value of a run. Second, we define the value of a word based on the values of its runs. To define values of runs, we will consider *value functions* f that assign real numbers to sequences of rationals. Given a non-empty word w , every run π of \mathcal{A} on w defines a sequence of weights of successive transitions of \mathcal{A} , i.e., $C(\pi) = (C(\pi[i-1], w[i], \pi[i]))_{1 \leq i \leq |w|}$; and the value $f(\pi)$ of the run π is defined as $f(C(\pi))$. We denote by $(C(\pi))[i]$ the weight of the i -th transition, i.e., $C(\pi[i-1], w[i], \pi[i])$. The value of a non-empty word w assigned by the automaton \mathcal{A} , denoted by $\mathcal{L}_{\mathcal{A}}(w)$, is the infimum of the set of values of all *accepting* runs; i.e., $\inf_{\pi \in \text{Acc}(w)} f(\pi)$, and we have the usual semantics that infimum of an empty set is infinite, i.e., the value of a word that has no accepting run is infinite. Every run π on an empty word has length 1 and the sequence $C(\pi)$ is empty, hence we define the value $f(\pi)$ as an external (not a real number) value \perp . Thus, the value of the empty word is either \perp , if the empty word is accepted by \mathcal{A} , or ∞ otherwise. To indicate a particular value function f that defines the semantics, we will call a weighted automaton \mathcal{A} an f -automaton.

Value functions. For finite runs we consider the following classical value functions: for runs of length $n + 1$ we have

- *Sum, absolute sum:* the sum function $\text{SUM}(\pi) = \sum_{i=1}^n (C(\pi))[i]$, the absolute sum $\text{SUM}^+(\pi) = \sum_{i=1}^n \text{Abs}((C(\pi))[i])$, where $\text{Abs}(x)$ is the absolute value of x ,

For infinite runs we consider:

- *Limit average:* $\text{LIMAVG}(\pi) = \liminf_{k \rightarrow \infty} \frac{1}{k} \cdot \sum_{i=1}^k (C(\pi))[i]$.

Silent moves. Consider a $(\mathbb{Z} \cup \{\perp\})$ -labeled automaton. We can consider such an automaton as an extension of a weighted automaton in which transitions labeled by \perp are *silent*, i.e., they do not contribute to the value of a run. Formally, for every function $f \in \text{InfVal}$ we define $\text{sil}(f)$ as the value function that applies f on sequences after removing \perp symbols. The significance of silent moves is as follows: it allows to ignore transitions, and thus provide robustness where properties could be specified based on desired events rather than steps.


2.2 Nested weighted automata

In this section we describe nested weighted automata introduced in [16], and closely follow the description of [16]. For more details and illustration of such automata we refer the reader to [16]. We start with an informal description.

Informal description. A *nested weighted automaton* (NWA) consists of a labeled automaton over infinite words, called the *master automaton*, a value function f for infinite words, and a set of weighted automata over finite words, called *slave automata*. A nested weighted automaton can be viewed as follows: given a word, we consider the run of the master automaton on the word, but the weight of each transition is determined by dynamically running slave automata; and then the value of a run is obtained using the value function f . That is, the master automaton proceeds on an input word as an usual automaton, except that before it takes a transition, it can start a slave automaton corresponding to the label of the current transition. The slave automaton starts at the current position of the word of the master automaton and works on some finite part of the input word. Once a slave automaton finishes, it returns its value to the master automaton, which treats the returned value as the weight of the current transition that is being executed. Note that for some transitions the master automaton might not invoke any slave automaton, and which corresponds to *silent* transitions. If one of slave automata rejects, the nested weighted automaton rejects. We first present an example and then the formal definition.

► **Example 1** (Average response time). Consider an alphabet Σ consisting of requests r , grants g , and null instructions $\#$. The average response time (ART) property asks for the average number of instructions between any request and the following grant. An NWA computing the average response time is depicted in Fig. 1. At every position with letter r the master automaton \mathcal{A}_{mas} of \mathbb{A} invokes the slave automaton \mathfrak{B}_1 , which computes the number of letters from its initial position to the first following grant. The automaton \mathfrak{B}_1 is a SUM^+ -automaton. On letters $\#$ and g the automaton \mathcal{A}_{mas} invokes the slave automaton \mathfrak{B}_2 , which is a dummy automaton, i.e., it immediately accepts and returns no weight. Invoking such a dummy automaton corresponds to taking a silent transition. Thus, the sequence of values returned by slave automata (54311... in Fig. 1), is the sequence of response times for each request. Therefore, the averages of these values is precisely the average response time; in the NWA \mathbb{A} , the value function f is LIMAVG. Also this property cannot be expressed by a non-nested automaton: a quantitative property is a function from words to reals, and as a function the range of non-nested LIMAVG-automata is bounded, whereas the ART can have unbounded values (for details see [16]).

Nested weighted automata. A *nested weighted automaton* (NWA) is a tuple $\langle \mathcal{A}_{mas}; f; \mathfrak{B}_1, \dots, \mathfrak{B}_l \rangle$, where (1) \mathcal{A}_{mas} , called the *master automaton*, is a $\{1, \dots, k\}$ -labeled



■ **Figure 1** An NWA \mathbb{A} computing ART. The master automaton \mathcal{A}_{mas} and slave automata $\mathfrak{B}_1, \mathfrak{B}_2$ are on the left. A part of a run of \mathbb{A} on word $rrr\#rgrg\dots$ is presented on the right.

automaton over infinite words (the labels are the indexes of automata $\mathfrak{B}_1, \dots, \mathfrak{B}_l$), (2) f is a value function on infinite words, called the *master value function*, and (3) $\mathfrak{B}_1, \dots, \mathfrak{B}_l$ are weighted automata over finite words called *slave automata*. Intuitively, an NWA can be regarded as an f -automaton whose weights are dynamically computed at every step by a corresponding slave automaton. We define an $(f; g)$ -automaton as an NWA where the master value function is f and all slave automata are g -automata.

Semantics: runs and values. A *run* of \mathbb{A} on an infinite word w is an infinite sequence $(\Pi, \pi_1, \pi_2, \dots)$ such that (1) Π is a run of \mathcal{A}_{mas} on w ; (2) for every $i > 0$ we have π_i is a run of the automaton $\mathfrak{B}_{C(\Pi[i-1], w[i], \Pi[i])}$, referenced by the label $C(\Pi[i-1], w[i], \Pi[i])$ of the master automaton, on some finite word of $w[i, j]$. The run $(\Pi, \pi_1, \pi_2, \dots)$ is *accepting* if all runs Π, π_1, π_2, \dots are accepting (i.e., Π satisfies its acceptance condition and each π_1, π_2, \dots ends in an accepting state) and infinitely many runs of slave automata have length greater than 1 (the master automaton takes infinitely many non-silent transitions). The value of the run $(\Pi, \pi_1, \pi_2, \dots)$ is defined as $\text{sil}(f)(v(\pi_1)v(\pi_2)\dots)$, where $v(\pi_i)$ is the value of the run π_i in the corresponding slave automaton. The value of a word w assigned by the automaton \mathbb{A} , denoted by $\mathcal{L}_{\mathbb{A}}(w)$, is the infimum of the set of values of all *accepting* runs. We require accepting runs to contain infinitely many non-silent transitions because f is a value function over infinite sequences, so we need the sequence $v(\pi_1)v(\pi_2)\dots$ with \perp symbols removed to be infinite.

Deterministic nested weighted automata. An NWA \mathbb{A} is *deterministic* if (1) the master automaton and all slave automata are deterministic, and (2) slave automata recognize prefix-free languages, i.e., languages \mathcal{L} such that if $w \in \mathcal{L}$, then no proper extension of w belongs to \mathcal{L} . Condition (2) implies that no accepting run of a slave automaton visits an accepting state twice. Intuitively, slave automata have to accept the first time they encounter an accepting state as they will not see an accepting state again.

► **Definition 2** (Width of NWA). An NWA has *width* k if and only if in every run at every position at most k slave automata are active.

► **Example 3** (Non-overlapping ART). Consider the NWA \mathbb{A} from Example 1 depicted in Fig. 1, which does not have bounded width. The run in Fig. 1 has width at least 4, but on word $rgr^2gr^3g\dots$ the number of active slave automata at position of letter g in subword $r^i g$ is i . We consider a variant of the ART property, called the 1-ART property, where after a request till it is granted additional requests are not considered. Formally, we consider the ART property over the language \mathcal{L}_1 defined by $(r\#^*g\#^*)^\omega$ (equivalently, given a request, the automata can check if the slave automaton is not active, and only then invoke it). An NWA

\mathbb{A}_1 computing the ART property over \mathcal{L}_1 is obtained from the NWA from Fig. 1 by taking the product of the master automaton \mathcal{A}_{mas} (from Fig. 1) with an automaton recognizing the language \mathcal{L}_1 . The automaton \mathbb{A}_1 , as well as, \mathbb{A} from Example 1 are (LIMAVG; SUM⁺)-automata and they are deterministic. Indeed, the master automaton and the slave automata of \mathbb{A}_1 (resp., \mathbb{A}) are deterministic and the slave automata recognize prefix-free languages. Moreover, in any (infinite) run of \mathbb{A}_1 at most one slave automaton is active, i.e., \mathbb{A}_1 has width 1. The dummy slave automata do not increase the width as they immediately accept, and hence they are not considered as active even at the position they are invoked. Finally, observe that the 1-ART property can return unbounded values, which implies that there exists no (non-nested) LIMAVG-automaton expressing it.

Decision problems. The classical questions in automata theory are language *emptiness* and *universality*. These problems have their counterparts in the quantitative setting of weighted automata and NWA. The (quantitative) emptiness and universality problems are defined in the same way for weighted automata and NWA; in the following definition the automaton \mathcal{A} can be either a weighted automaton or an NWA.

- **Emptiness:** Given an automaton \mathcal{A} and a threshold λ , decide whether there exists a word w with $\mathcal{L}_{\mathcal{A}}(w) \leq \lambda$.
- **Universality:** Given an automaton \mathcal{A} and a threshold λ , decide whether for every word w we have $\mathcal{L}_{\mathcal{A}}(w) \leq \lambda$.

The universality question asks for *non-existence* of a word w such that $\mathcal{L}_{\mathcal{A}}(w) > \lambda$.

► **Remark.** In this work we focus on value functions SUM and SUM⁺ for finite words, and LIMAVG for infinite words. There are other value functions for finite words, such as MAX, MIN and bounded sum. However, it was shown in [16] that for these value functions, there is a reduction to non-nested weighted automata. Also for infinite words, there are other value functions such as SUP, LIMSUP, where the complexity and decidability results have been established in [16]. Hence in this work we focus on the most conceptually interesting case of LIMAVG function for master automaton and the SUM and SUM⁺ value functions for the slave automata.

3 Examples

In this section we present several examples of properties of interest that can be specified with NWA of bounded width.

► **Example 4 (Variants of ART).** Recall the ART property (Example 1) and its variant 1-ART property (Example 3). We present two variants of the ART property.

First, we extend Example 3 and consider the k -ART property over languages \mathcal{L}_k defined by $(\#^* r (\#^* r \#^*)^{\leq k-1} g \#^*)^\omega$, i.e., the language where there are at most k -pending requests before each grant. As Example 3, an NWA \mathbb{A}_k computing the k -ART property can be constructed from the NWA from Fig. 1 by taking the product of the master automaton \mathcal{A}_{mas} (from Fig. 1) with an automaton recognizing \mathcal{L}_k . The NWA \mathbb{A}_k has width k .

Second, we consider the 1-ART[k] property, where $\Sigma = \{r_i, g_i : i \in \{1, \dots, k\}\} \cup \{\#\}$, i.e., there are k -different types of “request-grant” pairs. The 1-ART[k] property asks for the average number of instructions between any request and the following grant of the corresponding type. Moreover, we consider as for 1-ART property that for every i , between a request r_i and the following grant of the corresponding type g_i , there is no request r_i of the same type. The 1-ART[k] can be expressed with an (LIMAVG; SUM⁺)-automaton $\mathbb{A}_1^{[k]}$ of width bounded by k , which is similar to \mathbb{A}_1 from Example 3. Basically, the NWA $\mathbb{A}_1^{[k]}$

has k slave automata; for $i \in \{1, \dots, k\}$ the slave automaton \mathfrak{B}_i is invoked on letters r_i and it counts the number of steps to the following grant g_i . Additionally, the master automaton checks that for every i , between any two grants g_i , there is at most one request r_i .

In Examples 1, 3, and 4 we presented properties that can be expressed with (LIMAVG; SUM⁺)-automata. The following property of *average excess* can be expressed with slave automata with SUM value functions that have both positive and negative weights, i.e., it can be expressed by an (LIMAVG; SUM)-automaton, but not (LIMAVG; SUM⁺)-automata.

► **Example 5 (Block difference).** Consider the alphabet $\{r, g, \#\}$ from Example 1 with an additional letter $\$$. The *average excess* (AE) property asks for the average difference between requests and grants over blocks separated by $\$$. For example, for $\$(rr\#g\$)^\omega$ the average excess is 1. The AE property can be expressed by (LIMAVG; SUM)-automaton \mathbb{A}_{AE} of width 1 (presented below), but it cannot be expressed with (LIMAVG; SUM⁺)-automata; (LIMAVG; SUM⁺)-automata return values from the interval $[0, \infty)$, while AE ranges from $(-\infty, \infty)$. The automaton \mathbb{A}_{AE} invokes a slave automaton \mathfrak{B}_1 at positions of letter $\$$ and a dummy automaton \mathfrak{B}_2 on the remaining positions. The slave automaton \mathfrak{B}_1 runs until it sees $\$$ letter; it computes the difference between r and g letters by taking transitions of weights 1, $-1, 0$ respectively on letters $r, g, \#$. The master automaton as well as the slave automata of \mathbb{A}_{AE} are deterministic and the slave automata recognize prefix-free languages. Therefore, the NWA \mathbb{A}_{AE} is deterministic and has width 1.

4 Our Results

In this section we establish our main results. We first discuss complexity of checking whether a given NWA has width k . Next, we comment the results we need to prove. Afterwards, we present our results.

Configurations. Let \mathbb{A} be a non-deterministic (LIMAVG; SUM)-automaton of width k . We define a *configuration* of \mathbb{A} as a tuple $(q; q_1, \dots, q_k)$ where q is a state of the master automaton and each q_1, \dots, q_k is either a state of a slave automaton of \mathbb{A} or \perp . In the sequence q_1, \dots, q_k each state corresponds to one slave automaton, and the states are ordered w.r.t. the position when the corresponding slave automaton has been invoked, i.e., q_1 correspond to the least recently invoked slave automaton. If there are less than k slave automata active, then \perp symbols follow the actual states (denoting there is no slave automata invoked). We define $\text{CONF}(\mathbb{A})$ as the number of configurations of \mathbb{A} .

Key ideas. NWA without weights are equivalent to Büchi automata [16]. The property of having width k is independent from weights. It can be decided with a construction of a (non-weighted) Büchi automaton, which tracks configurations $(q; q_1, \dots, q_k)$ of a given NWA (assuming that it has width k) and accepts only if the width- k condition is at some point violated.

► **Theorem 6.** (1) Fix $k > 0$. We can check in polynomial time whether a given NWA has width k . (2) Given an NWA and a number k given in unary we can check in polynomial space whether the NWA has width k .

Proof. Let \mathbb{A} be an NWA and let $k > 0$ be a tested width. Consider a Büchi automaton \mathcal{A} , whose states are configurations of \mathbb{A} (considered to have width k) and a single accepting state q_{acc} . The automaton \mathcal{A} simulates runs of \mathbb{A} , i.e., it has a transition from one configuration to another over letter a if and only if there exist corresponding transitions of the master automaton and slave automata over letter a , which result in such a transition of \mathbb{A} . This

condition can be checked in polynomial time; for configurations $\langle q; q_1, \dots, q_k \rangle, \langle q'; q'_1, \dots, q'_k \rangle$ and letter a , we need to check whether $\langle q, a, q' \rangle$ is a transition of the master automaton of \mathbb{A} and each transition $\langle q_1, a, q'_1 \rangle, \dots, \langle q_k, a, q'_k \rangle$ is a transition of (some) slave automaton of \mathbb{A} . Additionally, whenever \mathcal{A} is in a state $\langle q; q_1, \dots, q_k \rangle$ such that $q_k \neq \perp$, i.e., k slave automata are active, and another slave automaton is invoked, \mathcal{A} takes a transition to q_{acc} , which is a single accepting state in \mathcal{A} . Observe that \mathcal{A} has an accepting run if and only if \mathbb{A} violates width- k condition.

The size of \mathcal{A} is bounded by $|\mathbb{A}|^k$, i.e., it is polynomial in the size of \mathbb{A} and exponential in k . Therefore, we can check emptiness of \mathcal{A} , and in turn violation of width- k condition, in polynomial time if k is constant and PSPACE if k is given in input in unary. ◀

Comment. We first note that for deterministic automata, emptiness and universality questions are similar. Hence we focus on the emptiness problem for non-deterministic automata (which subsumes the emptiness problem for deterministic automata) to establish the new results of Table 1. Moreover, the SUM^+ value function is a special case of the SUM value function with only positive weights. Since our main results are algorithms to establish upper bounds, we will only present the result for the emptiness problem for non-deterministic ($\text{LIMAVG}; \text{SUM}$)-automata. However, as a first step we show that without loss of generality, we can focus on the case of deterministic automata.

► **Lemma 7.** *Let $k > 0$. Given a non-deterministic ($\text{LIMAVG}; \text{SUM}$)-automaton \mathbb{A} over alphabet Σ of width k , a deterministic ($\text{LIMAVG}; \text{SUM}$)-automaton \mathbb{A}_d of width k over an alphabet $\Sigma \times \Gamma$ such that $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \inf_{u' \in (\Sigma \times \Gamma)^\omega} \mathbb{A}_d(u')$ can be constructed in time exponential in k and polynomial in $|\mathbb{A}|$. Moreover, $\text{CONF}(\mathbb{A}_d)$ is polynomial in $\text{CONF}(\mathbb{A})$ and k and only the alphabet of \mathbb{A}_d is exponential (in k) as compared to the alphabet of \mathbb{A} .*

► **Lemma 7.** *Let $k > 0$. Given a non-deterministic ($\text{LIMAVG}; \text{SUM}$)-automaton \mathbb{A} over alphabet Σ of width k , a deterministic ($\text{LIMAVG}; \text{SUM}$)-automaton \mathbb{A}_d of width k over an alphabet $\Sigma \times \Gamma$ such that $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \inf_{u' \in (\Sigma \times \Gamma)^\omega} \mathbb{A}_d(u')$ can be constructed in time exponential in k and polynomial in $|\mathbb{A}|$. Moreover, $\text{CONF}(\mathbb{A}_d)$ is polynomial in $\text{CONF}(\mathbb{A})$ and k and only the alphabet of \mathbb{A}_d is exponential (in k) as compared to the alphabet of \mathbb{A} .*

Proof. The general idea is to extend the alphabet to encode the actual letter and auxiliary symbols, which indicate how to resolve non-determinism. This can be done by explicitly writing down transitions the master and active slave automata should take. However, in such a solution, two slave automata which are at the same position in the same state, have the same suffixes of their runs, even though in a non-deterministic automaton their (suffixes) of runs can be different.

To circumvent this problem, we modify the automaton \mathbb{A} to \mathbb{A}' such that each slave automaton comes in k copies and the master automaton of \mathbb{A}' can invoke any copy of a slave automaton, i.e., if the master automaton of \mathbb{A} has a transition (q, a, q', i) , at which it invokes the slave automaton \mathfrak{B}_i , the master automaton of \mathbb{A} can invoke any copy of \mathfrak{B}_i . Clearly, \mathbb{A}' does not have any additional behaviors, i.e., by merging copies of slave automata in a run of \mathbb{A}' , we can obtain a run \mathbb{A} . Conversely, for every run of \mathbb{A} there exist multiple corresponding runs of \mathbb{A}' . In particular, for every run of \mathbb{A} there exists a run of \mathbb{A}' such that at every position, all active slave automata are different. We call such runs *controllable* as every slave automaton can be controlled independently of the others. In the following, we construct a deterministic automaton \mathbb{A}_d , which has corresponding run to every controllable run of \mathbb{A}' .

Before we describe the construction of \mathbb{A}_d , observe that without loss of generality, we can assume that accepting states of slave automata do not have outgoing transitions. Intuitively, we can clone each accepting state s into two copies s_1, s_2 , of which s_1 is a state with the transitions of s , but is not accepting, and s_2 is accepting but has no outgoing transitions. Then, all transitions to s are changed into two transitions, one to s_1 and one to s_2 .

Let \mathcal{Q} be the union of the set of states of the master automaton and all sets of states of slave automata of \mathbb{A}' . We define Γ as the set of partial functions h from $(k+1)$ -elements subsets of \mathcal{Q} into \mathcal{Q} . We define an $(\text{LIMAVG}; \text{SUM}^+)$ -automaton \mathbb{A}_d over the alphabet $\Sigma \times \Gamma$ by modifying only the transition relations and labeling functions of the master automaton and slave automata of \mathbb{A}' ; the sets of states and accepting states are the same as in the original automata. The transition relation and the labeling function of the master automaton \mathcal{A}_{mas}^d of \mathbb{A}_d is defined as follows: for all states q, q' , $(q, \langle a, h \rangle, q')$ iff $h(q) = q'$ and the master automaton of \mathbb{A}' has the transition (q, a, q') . The label of the transition $(q, \langle a, h \rangle, q')$ is the same as the label of the transition (q, a, q') . Similarly, for each slave automaton \mathfrak{B}_i in \mathbb{A}' , the transition relation of the corresponding slave automaton \mathfrak{B}_i^d in \mathbb{A}_d is defined as follows: for all states q, q' of \mathfrak{B}_i^d , $(q, \langle a, h \rangle, q')$ is a transition of \mathfrak{B}_i^d iff there $h(q) = q'$ and \mathfrak{B}_i has the transition (q, a, q') . The label of the transition $(q, \langle a, h \rangle, q')$ is the same as the label of the transition (q, a, q') in \mathfrak{B}_i^d . Observe that $(q, \langle a, h \rangle, q')$ can be a transition of \mathfrak{B}_i^d only if $h(q)$ is defined.

Observe that the master automaton of \mathbb{A}_d and all slave automata \mathfrak{B}_i^d are deterministic. Moreover, since we assumed that for every slave automaton in \mathbb{A}' final states have no outgoing transitions, slave automata \mathfrak{B}_i' recognize prefix free languages. Finally, it follows from the construction that (i) for every controllable run of \mathbb{A}' , there exists a corresponding run of \mathbb{A}_d where the sequence of values returned by slave automata is the same as in the run of \mathbb{A}' ; such runs have the same value. Basically, we encode in the input word, transitions of all automata with functions $h \in \Gamma$. Due to controllability of the run, at every position every slave automaton is in a different state. Therefore, we can encode in h transitions of all slave automata as well as the master automaton. Conversely, (ii) a run of \mathbb{A}_d , which is as a sequence of sequences of states, is basically a run of \mathbb{A} and it clearly has the same value. Therefore, the infimum over all words of \mathbb{A}_d coincides with the infimum over all words of \mathbb{A}' as well as of \mathbb{A} . \blacktriangleleft

Proof overview. We present our proof overview for the emptiness of deterministic $(\text{LIMAVG}; \text{SUM})$ -automata. The proof consists of the following four key steps.

1. First, we identify a condition, and show in Lemma 9 that it is a sufficient condition to ensure that the infimum value among all words is $-\infty$ (i.e., the least value possible). Moreover we show that the condition can be decided in PTIME if k is constant (even NLOGSPACE if additionally the weights are in unary) and in PSPACE if k is given in unary.
2. Second, we show that if the above condition does not hold, then there is a family of lasso words (i.e., a finite prefix followed by an infinite repetition of another finite word) that approximates the infimum value among all words. This shows that the above condition is both necessary and sufficient. Moreover, we consider *dense* words, where if we consider that slave automata have been invoked for the i -th time, then the run of the slave automata invoked is at most for $O(\log(i))$ steps. We show that the infimum is achieved by a dense word. These results are established in Lemma 11.
3. Third, we show using the above result, that the problem for bounded width can be reduced to the problem of width 1, and the reduction is polynomial in the size of the

original automaton, and only exponential in k . Thus if k is constant, the reduction is polynomial. This is established in Lemma 12.

4. Finally, we show that for automata with width 1, the emptiness problem can be solved in NLOGSPACE if weights are in unary and otherwise in PTIME (Lemma 14).

Given the above four steps we conclude our main result (Theorem 15). We start with the first item.

Intuition for the condition. We first illustrate with an example that for very similar automata, which just differ in order of invoking slave automata, the infimum over the values are very different. For one automaton the infimum value is $-\infty$ and for the other it is 0. This example provides the intuition for the need of the condition to identify when the infimum value is $-\infty$.


► **Example 8.** Consider two deterministic (LIMAVG; SUM)-automata $\mathbb{A}_1, \mathbb{A}_2$ defined as follows. The master automaton \mathcal{A}_{mas} of \mathbb{A}_1 accepts the language $(12a^*\#)^\omega$. At letter 1 (resp., 2) it invokes an automaton \mathfrak{B}_1 (resp., \mathfrak{B}_2). The slave automaton \mathfrak{B}_1 increments its value at every a letter and it terminates once it reads $\#$. The slave automaton \mathfrak{B}_2 works as \mathfrak{B}_1 except that it decrements its value at a letters. NWA \mathbb{A}_2 is similar to \mathbb{A}_1 except that it accepts the language $(21a^*\#)^\omega$. It invokes the same slave automata as \mathbb{A}_1 . Thus the two automata only differ in the order of invocation of the slave automata. Observe that the infimum over values of all words in \mathbb{A}_1 is 0. Basically, the values of slave automata are always the opposite, therefore the average of the values of slave automata is 0 infinitely often. However, the infimum over values of all words in \mathbb{A}_2 is $-\infty$. Indeed, consider a word $21a^1\# \dots 21a^{2^i} \dots$. At positions proceeding $1a^{2^i}$, the automaton \mathfrak{B}_2 returns the value -2^i and the average of all previous $2 \cdot i$ values is 0. Thus, the average at this position equals $-\frac{2^i}{2 \cdot i}$ (recall that the average is over the number of invocations of slave automata). Hence, the limit infimum of averages is $-\infty$.

Condition for infinite infimum. Let $k > 0$ and \mathbb{A} be a deterministic (LIMAVG; SUM)-automaton of width k . Let C be the minimal weight of slave automata of \mathbb{A} . Condition (*):
 (*) $C < 0$ and there exists a word w accepted by \mathbb{A} and infinitely many positions b such that the sum of weights, which automata active at position b accumulate while running on $w[b, \infty]$, is smaller than $C \cdot k^2 \cdot \text{CONF}(\mathbb{A})$.

Intuitively, condition (*) implies that there is a subword u which can be repeated so that the values of slave automata invoked before position b can be decreased arbitrarily. Note that pumping that word may not decrease the total average of the word. However, with LIMAVG value function, we need to ensure only the existence of a subsequence of positions at which the averages tend to $-\infty$, i.e., we only need to decrease the values of slave automata invoked before position b (for infinitely many positions).

Illustration of condition on example. Consider automata $\mathbb{A}_1, \mathbb{A}_2$ from Example 8. The automaton \mathbb{A}_2 satisfies condition (*), whereas \mathbb{A}_1 does not. In the word $21a^1\# \dots 21a^{2^i} \dots$, consider positions b , where \mathfrak{B}_2 is invoked by \mathbb{A}_2 . The automaton \mathfrak{B}_2 works on the subword $21a^{2^i}$, where both automata $\mathfrak{B}_1, \mathfrak{B}_2$ are active and the sum of their values past any position is 0. However, the only slave automaton active at position b is \mathfrak{B}_2 . These automaton accumulates the value -2^i past position b . Therefore, past some position N , all such positions b satisfy the statement from condition (*), and hence \mathbb{A}_2 satisfies condition (*). Now, for \mathbb{A}_1 , at every position at which \mathfrak{B}_2 is active, \mathfrak{B}_1 is active as well, hence for any position b , the values accumulated by slave automaton active past this position is non-negative. Hence, \mathbb{A}_1 does not satisfy condition (*). We now present our lemma about the condition.

► **Lemma 9.** Let $k > 0$ and \mathbb{A} be a deterministic (LIMAVG; SUM)-automaton of width k .



■ **Figure 2** Explanation to the proof in Lemma 9

- (1) If condition (*) holds for \mathbb{A} , then $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = -\infty$.
(2) Condition (*) can be checked in NLOGSPACE if the width is constant and weights are given in unary, PTIME if the width is constant, and in PSPACE if the width is given in unary.

Proof. We present proofs for each item below.

Proof of (1). Assume that (*) holds. We show that there exists a word u' such that $\mathbb{A}(u') = -\infty$. Consider a word w and position \mathbf{b} which (1) satisfy condition (*), and (2) the configuration at the position \mathbf{a} , the position of invocation of least recent slave automaton active at \mathbf{b} , occurs infinitely often in the run of \mathbb{A} in w .

At every step all slave automata decrease the sum of weights by at most $C \cdot k$. Therefore, there exist more than $k \cdot \text{CONF}(\mathbb{A})$ positions in $w[\mathbf{b}, \infty]$ at which the the sum of weights of automata invoked before position \mathbf{b} decreases. Thus there exist positions $\mathbf{x}_1, \mathbf{x}_2$ between which no automaton invoked before position \mathbf{b} terminates and there are more than $\text{CONF}(\mathbb{A})$ positions at which the sum of values of these automata decreases.

Therefore, there exist positions $\mathbf{c}_1, \mathbf{c}_2$ such that the configurations at $\mathbf{c}_1, \mathbf{c}_2$ are the same, and the sum of values of all automata activated before \mathbf{b} , which are still active at \mathbf{c}_2 decreases between \mathbf{c}_1 and \mathbf{c}_2 (see Fig. 2). Let \mathbf{d} be a position past \mathbf{c}_2 with the same configuration as at \mathbf{a} such that the master automaton visits an accepting state between \mathbf{a} and \mathbf{d} . Recall that all slave automata active at \mathbf{b} have been invoked past \mathbf{a} . We show how to build the word u' from subwords $w[1, \mathbf{a}]$, $w[\mathbf{a}, \mathbf{c}_1]$, $w[\mathbf{c}_1, \mathbf{c}_2]$ and $w[\mathbf{c}_2, \mathbf{d}]$.

Let $f(i) = 2^{2^i}$, i.e., $f(n+1) = f(n) \cdot f(n)$. Consider the word $u' = w[1, \mathbf{a}]\gamma_1\gamma_2\dots$, where $\gamma_i = w[\mathbf{a}, \mathbf{c}_1](w[\mathbf{c}_1, \mathbf{c}_2])^{f(i)}w[\mathbf{c}_2, \mathbf{d}]$. Observe that the sum of values of slave automata invoked in the prefix $w[1, \mathbf{a}]\gamma_1\dots\gamma_nw[\mathbf{a}, \mathbf{c}_1]$, which is shorter than $2 \cdot f(n)$ is less than $k \cdot C \cdot 2 \cdot f(n) - f(n+1) < -f(n)(f(n) - 2 \cdot k \cdot C)$. At most $2 \cdot f(n)$ slave automata have been invoked in this prefix, hence the average is less than $0.5 \cdot (-f(n) + k \cdot C)$. Hence, the limit infimum tends to minus infinity. Moreover, the master automaton visits one of its accepting states at least once in each γ_i .

Proof of (2). Consider a graph $G(\mathbb{A})$ of configurations of \mathbb{A} , in which there exists an edge from configuration CNF_1 to CNF_2 if and only if the automaton \mathbb{A} has a transition from CNF_1 to CNF_2 . Recall that in a configuration $(q; q_1, \dots, q_j, q_{j+1}, \dots, q_k)$ the states q_1, \dots, q_j correspond to the j least recently invoked slave automata. We show that condition (*) holds if and only if (**) there exists a cycle in $G(\mathbb{A})$ such that


- (a) a configuration CNF with an accepting state of the master automaton of \mathbb{A} is reachable from the cycle and the cycle is reachable from CNF , and
- (b) for some $j \geq 1$, the sum of weights of the j least recently invoked slave automata is negative in this cycle.
- (*) \Rightarrow (**): Consider word $w[1, \mathbf{a}]w[\mathbf{a}, \mathbf{c}_1]w[\mathbf{c}_1, \mathbf{c}_2]$. The configurations (in the run of \mathbb{A}) along the subword $w[\mathbf{c}_1, \mathbf{c}_2]$ define a cycle satisfying (a) and (b). Indeed, the master automaton of \mathbb{A} visits its accepting state on $w[1, \mathbf{a}]w[\mathbf{a}, \mathbf{c}_1]w[\mathbf{c}_1, \mathbf{c}_2]w[\mathbf{c}_2, \mathbf{d}]$; hence (a) holds. Let j be the number of slave automata invoked at position \mathbf{a} . Then, the sum of weights of the j least recently invoked slave automata on $w[\mathbf{c}_1, \mathbf{c}_2]$ is negative; hence (b) holds.
- (**) \Rightarrow (*): Let D be the maximal absolute weight in nested automata of \mathbb{A} . Consider a word which corresponds to a path in $G(\mathbb{A})$ from the initial configuration to the cycle satisfying (**) and repeating infinitely the following extended cycle: looping at the cycle from (**) $2 \cdot D \cdot k^2 \cdot \text{CONF}(\mathbb{A})$ times, and finally going through all configurations to terminate all slave automata, visit an accepting state of \mathbb{A} and returning to the start of the extended cycle. Observe that terminating all slave automata can be done in at most $k \cdot \text{CONF}(\mathbb{A})$ steps, therefore all slave automata active at the beginning of the extended cycle accumulate in total the weight smaller than $-2 \cdot D \cdot k^2 \cdot \text{CONF}(\mathbb{A})$ (iterating the negative cycle) plus $D \cdot k \cdot \text{CONF}(\mathbb{A})$ (terminating slave automata), which is smaller than $C \cdot k^2 \cdot \text{CONF}(\mathbb{A})$; hence (*) holds.

Finally, checking existence of such a cycle can be done in NLOGSPACE w.r.t. to the size of $G(\mathbb{A})$. If the width k is constant, then the size $G(\mathbb{A})$ is $O(|\mathbb{A}|^k)$, hence it is polynomial (in \mathbb{A}) if k is constant and exponential if k is given in unary. Thus, checking condition (*) is NLOGSPACE for constant width and PSPACE if the width is given in unary. \blacktriangleleft

► Definition 10. Let \mathbb{A} be a deterministic (LIMAVG; SUM)-automaton of width k . A word w is *dense* (w.r.t. \mathbb{A}) if in the run of \mathbb{A} on w , for every $i > 0$, the i -th invoked slave automaton takes at most $O(\log(i))$ steps.

Intuitive explanation of dense words. In a deterministic (LIMAVG; SUM)-automaton, the average is over the number of invoked slave automata, but in general, the returned values of the slave automata can be arbitrarily large as compared to the number of invocations, and hence the partial averages need not converge. Intuitively, in dense words, slave automata are invoked and terminated relatively densely, i.e., the length of their run depends on the number of slave automata invoked till this position. In consequence, the value they can accumulate is small w.r.t. the average, i.e., their absolute contribution to the sum of first n elements is $O(\log(n))$, and hence the contribution of the value a single slave automaton converges to 0 and partial averages converge on dense words.

Illustration on example. Consider an automaton \mathbb{A}_1 from Example 8. We discuss the definition of density on an example of word $w = 12a^1\#12a^3\#\dots12a^{2^{i+1}}\#\dots$, which is not dense (w.r.t. \mathbb{A}_1). Observe that at the position of subword $12a^{2^{i+1}}$, the partial average is 0. Once \mathfrak{B}_1 is invoked it returns value $2 \cdot i + 1$ and it is $(2 \cdot i + 1)$ -th invocation of a slave automaton. Hence, the average increases to 1 only to be decreased to 0 after invocation of \mathfrak{B}_2 . However, the word $w' = 12a^1\#(12a^2\#)^3\#\dots(12a^{2^{i+1}}\#)^{2^i}\#\dots$ is dense. Indeed, before the slave automata invoked at subword $12a^{2^{i+1}}\#$ there are at least $\sum_{j=1}^{i-1} 2^j = 2^i - 1$ invoked slave automata. Therefore, the value $2 \cdot i + 1$ returned by \mathfrak{B}_1 invoked on $12a^{2^{i+1}}$ changes the average by at most $\frac{2 \cdot i + 1}{2^i}$; as previously invoking \mathfrak{B}_2 in the next step bring the average back to 0. Therefore, the sequence of partial averages of values returned by slave automata converges to 0.



■ **Figure 3** Explanation to Lemma 11; the blue part corresponds to H , while the green part corresponds to T

► **Lemma 11.** *Let $k > 0$ and \mathbb{A} be a deterministic (LIMAVG; SUM)-automaton of width k . Assume that condition (*) does not hold. Then the following assertions hold:*

- (1) *For every $\epsilon > 0$ there exist finite words $\alpha_\epsilon, \beta_\epsilon$ such that $|\inf_{u \in \Sigma^\omega} \mathbb{A}(u) - \mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega)| < \epsilon$.*
- (2) *The value $\inf_{u \in \Sigma^\omega} \mathbb{A}(u)$ is greater than $-\infty$.*
- (3) *There exists a dense word w_d such that $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \mathbb{A}(w_d)$.*

Proof. We present the proof of each item.

Proof of (1): Consider $\epsilon > 0$. Let w_ϵ be a word such that $\mathbb{A}(w_\epsilon) - \inf_{u \in \Sigma^\omega} \mathbb{A}(u) < \frac{\epsilon}{4}$. There exists N_0 such that past N_0 , all configurations of \mathbb{A} occur infinitely often in the run of \mathbb{A} on w_ϵ and no position $b > N_0$ in w satisfies condition (*). Let $N_0 < a < b$ be positions in w_ϵ such that

- (1) the configuration of \mathbb{A} at a and b is the same,
- (2) the average of the values of slave automata invoked between a and b is at most $\mathbb{A}(w_\epsilon) + \frac{\epsilon}{4}$,
- (3) the number n , which is the number of slave automata invoked between positions a and b , is greater than $\frac{4}{\epsilon} \cdot C \cdot k^2 \cdot \text{CONF}(A)$, where C is the minimal weight, and
- (4) the sum of the weights slave automata active at a accumulate past a is at most $\frac{n \cdot \epsilon}{4}$,
- (5) all slave automata invoked before a are terminated before b ,
- (6) the master automaton of \mathbb{A} visits between a and b an accepting state at least once.

Let $P > N_0$ be the minimal position such that all configurations, which appear infinitely often the run of \mathbb{A} on w_ϵ , appear between N_0 and P . Let M be the maximal absolute sum of slave automata invoked before P . We pick b such that the average of all slave automata invoked before b is at most $\mathbb{A}(w_\epsilon) + \frac{\epsilon}{8}$, and b is large enough to satisfy (3), (6) and (4) and (5) for $a = P$. For such b we choose a from the interval $[N_0, P]$ and observe that (1)–(6) are satisfied.

We put $\alpha_\epsilon = u[1, a]$ and $\beta_\epsilon = u[a, b]$. We claim that $|\mathbb{A}(w_\epsilon) - \mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega)| < \epsilon$. Clearly, $\mathbb{A}(w_\epsilon) < \mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega)$. Conversely, due to condition (3), we have $\mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega) < \mathbb{A}(w_\epsilon) + \frac{\epsilon}{4} - T + H$, where T , the tail, is the sum of weights, which automata invoked before position b accumulate past position b , and H , the head, is the sum of weights, which automata invoked before position b accumulate while running on β_ϵ again (see Fig 3). By the condition (5), we can estimate H by $\frac{n \cdot \epsilon}{4}$. Due to (4) and by condition (*), we have $T < C \cdot k^2 \cdot \text{CONF}(\mathbb{A})$.

Proof of (2): Assume that condition (*) does not hold. Then, for $\epsilon = 1$ there exist $\alpha_\epsilon, \beta_\epsilon$ such that $|\mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega) - \inf_{u \in \Sigma^\omega} \mathbb{A}(u)| < 1$. However, observe that $\mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega)$ is finite. Indeed, all slave automata invoked in β_ϵ has to terminate within $|\beta_\epsilon|$ steps, otherwise there is a slave automaton with an infinite run. It follows that the value of $\mathbb{A}(\alpha_\epsilon(\beta_\epsilon)^\omega)$ is greater or equal to $C \times k \times |\beta_\epsilon|$, where C is the minimal weight in all slave automata of \mathbb{A} . In consequence, $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) > -\infty$.

Proof of (3): We can strengthen (1) and say that there exist pairs of words

$(v_1, u_1), (v_2, u_2), \dots$ such that (1) for every i , we have $|\mathbb{A}(v_i u_i^\omega) - \mathbb{A}(w_\epsilon)| < \frac{1}{i}$, and (2) for all i, j , the automaton \mathbb{A} reaches the same configuration from the initial state upon reading v_i and v_j . Having conditions (1) and (2) we define word w as $\alpha_1 \beta_1^{k[1]} \beta_2^{k[2]} \dots$, where the sequence $k[i]$ is defined as follows: $k[1] = 1$ and for $i \geq 1$ we have $k[i] = 2^{|\beta_{i-1}| + |\beta_i| + |\beta_{i+1}|}$. Clearly, $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \mathbb{A}(w')$. Observe that every slave automaton invoked at a subword β_i terminates within $\max(|\beta_i|, |\beta_{i+1}|)$ steps, depending on whether the following subword in β_i or β_{i+1} . However, the first subword β_i appears first after at least $2^{|\beta_i| + |\beta_{i+1}|}$ slave automata has been invoked. Therefore, the word w' is dense. \blacktriangleleft

► **Remark.** Lemma 9 together with (2) of Lemma 11 imply that for a deterministic (LIMAVG; SUM)-automaton \mathbb{A} of width k condition (*) is both necessary and sufficient for the infimum over all values equal to $-\infty$. Moreover, this condition can be checked efficiently.

Lemma 12 reduces the emptiness problem for deterministic (LIMAVG; SUM)-automata of width k to the same problem with automata of width 1.

► **Lemma 12.** *Let $k > 0$ and \mathbb{A} be a deterministic (LIMAVG; SUM)-automaton of width k . Assume that condition (*) does not hold. Then, there exists a deterministic (LIMAVG; SUM)-automaton \mathbb{A}_1 of width 1 over an alphabet Δ such that $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \inf_{u \in \Delta^\omega} \mathbb{A}_1(u)$. The size of \mathbb{A}_1 is $O(|\mathbb{A}|^k)$ and it can be constructed on-the-fly.*

Proof. We define the automaton \mathbb{A}_1 of width bounded by 1, whose slave automaton \mathfrak{B}^Σ simulates runs of k slave automata of \mathbb{A} . The slave automaton \mathfrak{B}^Σ computes the sum of weights collected by all active slave automata of \mathbb{A} . Once \mathbb{A} invokes a new slave automaton, \mathfrak{B}^Σ terminates and it is immediately restarted in the next transition. Moreover, if the last active slave automaton of \mathbb{A} terminates, \mathfrak{B}^Σ terminates as well. More formally, the master automaton of \mathbb{A}_1 is $(\Sigma, Q, q_0, \delta, F, C)$, where $Q = Q_m \times (Q_s \times \dots \times Q_s) \times \{0, 1\} = Q_m \times Q_s^k \times \{0, 1\}$, where Q_m is the set of states of the master automaton of \mathbb{A} and Q_s is the union of the set of states of The master automaton keeps track of the master automaton of \mathbb{A} and all its active slave automata. Its last bit $\{0, 1\}$ indicates whether the slave automaton is active, i.e., if one of slave automata of \mathbb{A} terminates, the bit is flipped to 0, and if the bit is 0 but some of slave automata of \mathbb{A} are active, then start the slave automaton in the configuration corresponding to the current configuration of slave automata. In the generalized Büchi acceptance condition F , we encode that the run of the master automaton accepts and runs of all slave automata are finite. The components q_0, δ, F, C are defined accordingly to the description.

The slave automaton \mathfrak{B}^Σ has the similar structure to the master automaton of \mathbb{A}_1 ; there are two key difference. First, if any of tracked slave automata of \mathbb{A} terminates, \mathfrak{B}^Σ terminates as well. Second, for every transition of \mathfrak{B}^Σ , the weight of this transition is the sum of weights of current transitions of tracked slave automata.

Now, we show that on dense words (w.r.t. \mathbb{A}), values of both automata coincide. Let w be a dense word. Consider a position i in w . Let n be the number of slave automata invoked by \mathbb{A} before position i . Observe that \mathbb{A}_1 invokes a new slave automaton whenever \mathbb{A} does. Therefore, \mathbb{A}_1 also invoked n slave automata before position i . The partial average up to position i in $\mathbb{A}(w)$ is the sum of values of all slave automata invoked before position i , while in $\mathbb{A}_1(w)$ this is the sum of values all slave automata invoked before position i accumulate before position $j > i$, the first position past i when a new slave automaton is invoked. Therefore, the difference between partial averages of $\mathbb{A}(w)$ and $\mathbb{A}_1(w)$ up to position i , denoted by Δ , is the sum of weights slave automata invoked before the position i accumulate past position j multiplied by $\frac{1}{n}$. Now, due to density of word w , each slave automaton invoked before the position i works for at most $\log(n)$ steps, therefore the absolute accumulated value is at

most $C \log(n)$, where C is the maximal absolute weight in the slave automata of \mathbb{A} . Hence, $|\Delta| < \frac{1}{n} \cdot k \cdot C \cdot \log(n)$. Therefore, the partial averages of $\mathbb{A}(w)$ and $\mathbb{A}_1(w)$ converge. In consequence, $\mathbb{A}(w) = \mathbb{A}_1(w)$.

Due to Lemma 11, there exist a dense word w (w.r.t. \mathbb{A}), which has the minimal value among all words. Then, we have $\mathbb{A}(w) = \mathbb{A}_1(w)$. Therefore, $\inf_u \mathbb{A}_1(u) \leq \inf_u \mathbb{A}(u)$. Conversely, we show that there exists a *optimal* word w such that $\inf_u \mathbb{A}_1(u) = \mathbb{A}_1(w)$ and w is a dense word (w.r.t. \mathbb{A}). Notice that every time \mathbb{A}_1 encodes in its generalized Büchi condition that all slave automata terminate infinitely often. It follows that between two positions a, b at which the generalized Büchi condition is satisfied, all slave automata invoked before a terminate before b . Thus, to show that w is a dense word (w.r.t. \mathbb{A}), we show that the generalized Büchi condition of \mathbb{A}_1 is satisfied at w sufficiently densely. Observe that a deterministic (LIMAVG; SUM)-automaton of width bounded by 1 is essentially equivalent to a LIMAVG-automaton with ϵ -transitions (see Lemma 14), for which one can construct words of optimal values which satisfies their (generalized) Büchi condition arbitrarily densely as long as the density converges to 0. Thus, such a word w exists, $\mathbb{A}(w) = \mathbb{A}_1(w)$. It follows that $\inf_u \mathbb{A}_1(u) = \inf_u \mathbb{A}(u)$. ◀

► **Claim 13.** *The emptiness problem for deterministic LIMAVG-automata with weights given in unary is NLOGSPACE-complete.*

Proof. NLOGSPACE-hardness readily follows from NLOGSPACE-hardness of directed graph reachability.

For containment in NLOGSPACE, recall that the infimum over all words of the values of a given LIMAVG-automaton \mathcal{A} is less or equal to a given λ if and only if there exists a cycle in the automaton such that

- (1) the cycle is reachable from the initial state,
- (2) the sum of weights along the cycle is less or equal to λ , and
- (3) there exists an accepting state s_a and a state s_c such that s_a is reachable from s_c and vice versa.

To check conditions (1),(2), (3) we non-deterministically pick states s_a, s_c and verify conditions (1) and (3) with reachability queries, which are in NLOGSPACE. To check (2) we observe that if there exists a cycle satisfying conditions (1),(2), (3), then there also exists a cycle that satisfy these conditions and of length bounded by $C \cdot |\mathcal{A}|$, where C is the maximal absolute value of the weights from \mathcal{A} . Now, if weights in \mathcal{A} are given in unary, $C < |\mathcal{A}|$, the length of the cycle is at most $|\mathcal{A}|^2$ and we need only logarithmic memory to non-deterministically pick the cycle state by state and keep track of the sum to verify that it is less or equal to λ . ◀

► **Lemma 14.** *The emptiness problem for deterministic (LIMAVG; SUM)-automata of width 1 is in PTIME and if the weights are in unary, then it is in NLOGSPACE.*

Proof. We observe that \mathbb{A} is essentially a deterministic LIMAVG-automaton with silent moves. More precisely, let Q_m be the set of states of the master automaton of \mathbb{A} and let n be the number of slave automata. A run of the automaton of width bounded by 1, can be partitioned into two types of fragments:

- (1) fragments corresponding to a single run of slave automata, which are characterized by $q_1, q_2 \in Q_m$, the state of the master automaton at beginning and at the end of the fragment, and
- (2) fragments where no slave automaton is running, i.e., only dummy slave automata are invoked, which are characterized by:

- $q_1, q_2 \in Q_m$, the state of the master automaton at beginning and at the end of the fragment,
- the first letter of the fragment a ,
- the index $i \in \{1, \dots, n\}$ of the slave automaton invoked, and
- the value returned by the invoked slave automaton.

We consider a succinct representation of runs of \mathbb{A} , where fragments of type (1) are replaced by a single letter (q_1, q_2) and fragments of type (2) are substituted by a single letter (q_1, a, q_2, i) . Moreover, we consider only maximal fragments, i.e., two fragments of type (1) can be merged into one fragment, therefore we forbid two successive occurrences of letters $(q_1, q_2)(q_2, q_3)$.

Let $\Delta = \{(q_1, q_2) : q_1, q_2 \in Q_m\} \cup \{(q_1, a, q_2, i) : q_1, q_2 \in Q_m, a \in \Sigma, i \in \{1, \dots, n\}\}$. We can define a deterministic LIMAVG-automaton \mathcal{A} with silent moves over Δ which accepts only words that represent accepting runs of \mathbb{A} . The automaton \mathcal{A} checks that for two successive letters a, b the second state in a is the same as the first state in b (e.g., $(q_1, q_2, i)(q_2, q_3)(q_3, q_4, i')$), and it has a list of valid letters, i.e., a letter are valid if there exists a fragment corresponding to it. More precisely, (q_1, q_2, i) is valid iff there exists a word v such that that (1) the master automaton in the state q_1 upon reading letter $v[1]$ takes a transition at which it invokes \mathfrak{B}_i , (2) the master automaton moves from q_1 to q_2 upon reading v , and (3) slave automaton \mathfrak{B}_i accepts the word v . Validity of letters (q_1, q_2) is defined similarly.

Moreover, transitions over letters (q_1, q_2) are silent (have no value) and transitions over letters (q_1, q_2, i) have the minimal value associated with such a fragment, i.e., it is the minimal value slave automaton \mathfrak{B}_i can return on a word v such that the master automaton moves on v from state q_1 to q_2 . Thus, the value of \mathcal{A} on $w' \in \Delta$ is (provided it is accepted) the minimal value \mathbb{A} can return on the run with the sequence of fragments corresponding to w' . Every run has the corresponding sequence of fragments In consequence, $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \inf_{u' \in \Delta^\omega} \mathbb{A}(u')$.

Finally, observe that \mathcal{A} can be constructed upon demand from \mathbb{A} in logarithmic space, i.e., we can answer query about parts of \mathcal{A} without outputting the whole automaton. The emptiness problem for LIMAVG-automata is decidable in NLOGSPACE if weights are given in unary notation and in PTIME if weights are given in binary (Claim 13). The silent moves are interleaved with non-silent moves, therefore we can easily remove them and decide the emptiness problem for \mathcal{A} , and in turn the emptiness problem for \mathbb{A} , (a) in NLOGSPACE for unary weights, and (b) in PTIME for binary weights. ◀

Key intuitions. We show that every transition of \mathcal{A}_{mas} , the master automaton of \mathbb{A} , at which a slave automaton is invoked, can be substituted by a transition whose weight is the minimal value the invoked slave automaton can achieve. More precisely, while a slave automaton is running on the input word, the master automaton \mathcal{A}_{mas} is still active. Therefore, we substitute transitions (q, a, q', i) of \mathcal{A}_{mas} by multiple transitions of the form $(q, (q, a, i, q''), q'')$, where (q, a, i, q'') is a new letter, q'' is a state of \mathcal{A}_{mas} and the weight of this transition is the minimal value \mathfrak{B}_i can achieve over words au such that \mathcal{A}_{mas} moves from q to q'' upon reading au . Such a transformation preserves the infimum over all words and it transforms a deterministic (LIMAVG; SUM)-automaton of width 1 to a deterministic LIMAVG-automaton. The emptiness problem for LIMAVG-automaton is decidable in PTIME and even in NLOGSPACE provided that weights are given in unary.

We now present the algorithm and lower bound for our main result.

The algorithm. We present an algorithm, which, given a non-deterministic (LIMAVG, SUM)-automaton \mathbb{A} of width k and $\lambda \in \mathbb{Q}$, decides whether $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) \leq \lambda$.

1. Transform \mathbb{A} into a deterministic (LIMAVG, SUM)-automaton \mathbb{A}_d of the same width such that $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = \inf_{u \in (\Sigma \times \Gamma)^\omega} \mathbb{A}_d(u)$ (Lemma 7).
2. Check condition (*) for \mathbb{A}_d . If it holds, then $\inf_{u \in \Sigma^\omega} \mathbb{A}(u) = -\infty$ and return answer **YES**. Otherwise, continue the algorithm.
3. Transform \mathbb{A}_d into a deterministic (LIMAVG, SUM)-automaton \mathbb{A}_1 of width 1 such that $\inf_{u \in (\Sigma \times \Gamma)^\omega} \mathbb{A}_d(u) = \inf_{u \in \Delta^\omega} \mathbb{A}_1(u)$ (Lemma 12).
4. Compute $\inf_{u \in \Delta^\omega} \mathbb{A}_1(u)$ (Lemma 14), and return whether $\inf_{u \in \Delta^\omega} \mathbb{A}_1(u) \leq \lambda$.

Transformations in (1) and (3) are polynomial in the size of the automaton and exponential in k . Also, transformation from (1) does not increase k . Therefore, the size of \mathbb{A}_1 is polynomial in the size \mathbb{A} and singly exponential in k . Moreover, these transformations can be done on-the-fly, i.e., there is not need to store the whole resulting automaton. Therefore, checks from (2) and (4), can be done in NLOGSPACE if k is constant and weights are in unary, PTIME if k is constant, and PSPACE if k is given in unary.

Hardness results. If k is constant, then the reachability problem on directed graphs, which is NLOGSPACE-complete, can be reduced to language emptiness of a finite automaton, which is a special case the emptiness problem for non-deterministic (LIMAVG, SUM)-automata of width 1 with unary weights. If k is given in unary, consider the emptiness problem for the intersection of regular languages, which given k and regular languages $\mathcal{L}_1, \dots, \mathcal{L}_k$, asks whether $\mathcal{L}_1 \cap \dots \cap \mathcal{L}_k = \emptyset$. This problem is PSPACE-complete [25] and reduces to the emptiness problem for deterministic (LIMAVG, SUM)-automata of width given in unary: the PSPACE-hardness result for emptiness of NWA given in [16] uses NWA of width $|\mathbb{A}|$.

► **Theorem 15.** *The emptiness problem for non-deterministic (LIMAVG, SUM)-automata is (a) NLOGSPACE-complete in the size of \mathbb{A} for constant width k with weights in unary; (b) PTIME in the size of \mathbb{A} for constant width k ; and (c) PSPACE-complete when the bounded width k is given as input in unary.*

Acknowledgements. This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE/SHiNE) and Z211-N23 (Wittgenstein Award), ERC Start grant (279307: Graph Games), Vienna Science and Technology Fund (WWTF) through project ICT15-003 and by the National Science Centre (NCN), Poland under grant 2014/15/D/ST6/04543.

References

- 1 Shaull Almagor, Udi Boker, and Orna Kupferman. Discounting in LTL. In *TACAS, 2014*, pages 424–439, 2014.
- 2 Rajeev Alur, Loris D’Antoni, Jyotirmoy V. Deshmukh, Mukund Raghothaman, and Yifei Yuan. Regular functions and cost register automata. In *LICS 2013*, pages 13–22, 2013.
- 3 Christel Baier, Clemens Dubslaff, and Sascha Klüppelholz. Trade-off analysis meets probabilistic model checking. In *CSL-LICS 2014*, pages 1:1–1:10, 2014.
- 4 Christel Baier, Joachim Klein, Sascha Klüppelholz, and Sascha Wunderlich. Weight monitoring with linear temporal logic: complexity and decidability. In *CSL-LICS 2014*, pages 11:1–11:10, 2014.
- 5 Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orna Kupferman. Temporal specifications with accumulative values. *ACM TOCL*, 15(4):27:1–27:25, 2014.
- 6 Benedikt Bollig, Paul Gastin, Benjamin Monmege, and Marc Zeitoun. Pebble weighted automata and transitive closure logics. In *ICALP 2010, Part II*, pages 587–598. Springer, 2010.

- 7 Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In *CONCUR 2014*, pages 266–280, 2014.
- 8 Tomás Brázdil, Václav Brozek, Krishnendu Chatterjee, Vojtech Forejt, and Antonín Kucera. Two views on multiple mean-payoff objectives in Markov decision processes. In *LICS 2011*, pages 33–42, 2011.
- 9 Tomás Brázdil, Krishnendu Chatterjee, Vojtech Forejt, and Antonín Kucera. Multigain: A controller synthesis tool for MDPs with multiple mean-payoff objectives. In *TACAS 2015*, pages 181–187, 2015.
- 10 Krishnendu Chatterjee. Markov decision processes with multiple long-run average objectives. In *FSTTCS*, pages 473–484, 2007.
- 11 Krishnendu Chatterjee and Laurent Doyen. Energy and mean-payoff parity Markov Decision Processes. In *MFCSS 2011*, pages 206–218, 2011.
- 12 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Alternating weighted automata. In *FCT’09*, pages 3–13. Springer, 2009.
- 13 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Expressiveness and closure properties for quantitative languages. *LMCS*, 6(3), 2010.
- 14 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM TOCL*, 11(4):23, 2010.
- 15 Krishnendu Chatterjee, Vojtech Forejt, and Dominik Wojtczak. Multi-objective discounted reward verification in graphs and MDPs. In *LPAR*, pages 228–242, 2013.
- 16 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Nested weighted automata. In *LICS 2015*, pages 725–737, 2015.
- 17 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Quantitative automata under probabilistic semantics. *CoRR*, abs/1604.06764, 2016. A conference version accepted to LICS 2016.
- 18 Krishnendu Chatterjee, Zuzana Komárková, and Jan Kretínský. Unifying two views on multiple mean-payoff objectives in Markov Decision Processes. In *LICS 2015*, pages 244–256, 2015.
- 19 Krishnendu Chatterjee, Rupak Majumdar, and Thomas A. Henzinger. Markov Decision Processes with multiple objectives. In *STACS 2006*, pages 325–336, 2006.
- 20 Manfred Droste, Werner Kuich, and Heiko Vogler. *Handbook of Weighted Automata*. Springer, 1st edition, 2009.
- 21 Manfred Droste and George Rahonis. Weighted automata and weighted logics on infinite words. In *DLT 2006*, pages 49–58, 2006.
- 22 Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer, 1996.
- 23 Vojtech Forejt, Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Hongyang Qu. Quantitative multi-objective verification for probabilistic systems. In *TACAS*, pages 112–127, 2011.
- 24 Thomas A. Henzinger and Jan Otop. From model checking to model measuring. In *CONCUR 2013*, pages 273–287, 2013.
- 25 Dexter Kozen. Lower bounds for natural proof systems. In *FOCS*, pages 254–266. IEEE Computer Society, 1977.
- 26 Mehryar Mohri. Semiring frameworks and algorithms for shortest-distance problems. *J. Aut. Lang. & Comb.*, 7(3):321–350, 2002.
- 27 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1st edition, 1994.